

Deciding reducibility of mapping classes is in **NP**

Mark C. Bell
University of Illinois
mcbell@illinois.edu

October 27, 2015

Abstract

For a fixed marked surface S , we show that the problem of deciding whether or not a mapping class is reducible lies in **NP**. As usual this immediately gives an exponential time algorithm to decide whether or not a mapping class is reducible.

To do this we use an (ideal) triangulation to obtain a coordinate system on the set of multicurves on S . The result then follows from the fact that the action of the mapping class group of S is piecewise-linear with respect to such a coordinate system and so we are able to show that: if a mapping class h fixes a multicurve then it fixes one whose size is at most exponential in the word length of h .

We go on to show how to repeat this construction on invariant sub-surfaces. This allows us to show that a similar bound holds for the size of the canonical curve system of a mapping class and so give an alternate, elementary proof of a result of Koberda and Mangahas.

keywords. mapping class group; reducible; canonical curve system; **NP**.

Mathematics Subject Classification (2010): 57M99

1 Introduction

Fix S to be a (possibly disconnected) marked surface in which each component contains at least one marked point and no component is a once or twice marked sphere.

Let $\mathcal{C}(S)$ denote the set of essential, simple, closed *multicurves* on S . This is strictly larger than the set of simplices of the curve complex of S [8]; it includes multicurves in which some of the components are parallel.

Let $\text{Mod}^+(S)$ denote the *mapping class group* of S , relative to the set of marked points. We fix X to be a finite generating set of $\text{Mod}^+(S)$ and let X^* denote the set of all words that can be made using the elements of X as letters. We identify a word $h = h_1 \cdots h_k \in X^*$ with the mapping class

$$h_k \circ \cdots \circ h_1$$

and denote its *length* by $\ell(h) := k$.

Definition 1.1. A mapping class $h \in \text{Mod}^+(S)$ is *reducible* if there is an h -*invariant* multicurve, that is, a multicurve $\gamma \in \mathcal{C}(S)$ such that $h(\gamma) = \gamma$. A word is *reducible* if its corresponding mapping class is.

Problem 1.2 (The Reducibility Problem). Given a word $h \in X^*$, decide whether or not h is reducible.

We begin in Section 2 by describing how to obtain a coordinate system on $\mathcal{C}(S)$ from an (ideal) triangulation. This allows us to efficiently represent multicurves by a vector of integers. In Section 2.1 we also use these triangulations to give a combinatorial way of representing mapping classes and in Section 2.2 we show how this can be used to efficiently compute the image of a multicurve under a mapping class.

In Section 3 we use the fact that the action of $\text{Mod}^+(S)$ on $\mathcal{C}(S)$ is piecewise-linear with respect to this coordinate system to show that: if there is an h -invariant multicurve then there is one that is small with respect to our coordinate system (Corollary 3.4). Such an invariant multicurve acts as a certificate of reducibility and is sufficiently small that it can be verified that it is invariant in polynomial time. Thus we deduce that:

Corollary 3.5. *The reducibility problem is in NP.*

Finally, in Section 4, we describe how to repeat this construction on invariant subsurfaces. Using these structures again, we obtain similar bounds and so are able to also give an elementary proof of a result of Koberda and Mangahas [7, Theorem 1]:

Corollary 4.9. *Fix \mathcal{T} , a triangulation of S . For each word $h \in X^*$, the edge vector $\mathcal{T}(\sigma(h))$ of the canonical curve system of h is $O(\ell(h))$ -bounded.*

Unlike the proof given by Koberda and Mangahas, the proof of this bound does not rely on any knowledge of the proof of solvability of the conjugacy problem for $\text{Mod}^+(S)$, constants related to the curve complex or the finite index classifying covers of S .

1.1 Notation

We begin by setting some general notation and noting that all vectors and matrices will have integer entries throughout. Let:

- \mathbb{N} denote the set of natural numbers including zero,
- $|S|$ denote the number of components of S ,
- $g(S)$ denote the genus of S , which we define to be the sum of the genres of the components when S is disconnected,
- $n(S)$ denote the number of marked points on S ,
- $v[i]$ denote the i^{th} entry of a vector v ,
- $v \geq 0$ denote that the vector v is *non-negative*, that is, each entry of v is non-negative,
- $v \geq_2 0$ denote that $v \geq 0$ and that each entry of v is even, and
- $\begin{pmatrix} A \\ B \end{pmatrix}$ denote the join of matrices A and B , obtained by stacking their rows.

1.2 Model of computation

To simplify our analysis we will assume that each variable can hold an arbitrary integer and there is no cost associated to variable access. However, in the problems that we will tackle not only will matrices grow in size with the problem but so will the entries involved. This growth is almost always exponential and so we must take extra care in our analysis.

Definition 1.3. An integer, vector or matrix is k -bounded if the log of the absolute value of each entry is at most k . That is, if each number involved can be represented by at most k bits.

As part of our model of computation, we will assume that if x and y are k -bounded and k' -bounded integers respectively where $k \geq k'$ then:

- $\text{sign}(x)$ can be computed in $O(1)$ operations,
- $x \pm y$ is $(k + 1)$ -bounded and can be computed in $O(k)$ operations, and
- xy is $(k + k')$ -bounded and can be computed in $O(kk')$ operations.

2 Triangulations

Definition 2.1. An (*ideal*) *triangulation* \mathcal{T} of S is the isotopy class of a finite, ordered collection of arcs on S which connect between the marked points, have pairwise disjoint interiors and are such that the metric completion of each component of $S - \mathcal{T}$ is an unmarked triangle.

When working with a triangulation, we refer to the marked points as *vertices*, the arcs as *edges* and the metric completion of each component of $S - \mathcal{T}$ as *faces*. We let

$$\zeta = \zeta(S) := 6g(S) + 3n(S) - 6|S|$$

denote the *complexity* of S . This is the number of edges of any triangulation of S .

The fact that the edges of a triangulation are ordered will be crucial. Changing the ordering of the edges of a triangulation does not produce an equivalent triangulation.

Definition 2.2. Let \mathcal{T} be a triangulation of S with edges e_1, \dots, e_ζ (in order). The *edge vector* of a multicurve $\gamma \in \mathcal{C}(S)$ with respect to \mathcal{T} is the vector

$$\mathcal{T}(\gamma) := \begin{pmatrix} \iota(\gamma, e_1) \\ \vdots \\ \iota(\gamma, e_\zeta) \end{pmatrix} \in \mathbb{N}^\zeta$$

where $\iota(x, y)$ is the *geometric intersection number* of x and y .

Although for each triangulation \mathcal{T} the map $\mathcal{T}(\cdot): \mathcal{C}(S) \rightarrow \mathbb{N}^\zeta$ is injective it is not surjective. In fact a vector $v \in \mathbb{N}^\zeta$ corresponds to a multicurve if and only if:

- for each face of \mathcal{T} with edges e_i, e_j and e_k we have that $v[i] + v[j] - v[k] \in 2\mathbb{N}$, and

- for each vertex v of \mathcal{T} there is a face with edges e_i , e_j and e_k such that $e_i \cap e_j = v$ and $v[i] + v[j] - v[k] = 0$.

We express this requirement as a linear programming problem:

Lemma 2.3. *For each triangulation \mathcal{T} , there are $O(1)$ -bounded $\zeta \times 3\zeta$ matrices F_1, \dots, F_k such that $v \in \mathbb{N}^\zeta$ is in the image of $\mathcal{T}(\cdot)$ if and only if $v \neq 0$ and*

$$F_i \cdot v \geq_2 0$$

for some i . □

This lemma allows us to test whether a k -bounded vector is in the image of $\mathcal{T}(\cdot)$ in $O(k)$ operations.

2.1 The graph of (ordered) triangulations

The edge vector of a multicurve depends heavily on the choice of triangulation. We consider two elementary ways of altering a triangulation, both of which change the edge vector of a multicurve predictably.

Firstly, we may use a permutation to reorder the edges of a triangulation.

Lemma 2.4. *Let \mathcal{T} and \mathcal{T}' be triangulations of S which are equivalent up to reordering. Let Σ be the permutation matrix corresponding to the reordering. Then for each multicurve $\gamma \in \mathcal{C}(S)$,*

$$\mathcal{T}'(\gamma) = \Sigma \cdot \mathcal{T}(\gamma). \quad \square$$

Secondly, if the interior of an edge e meets two distinct faces of \mathcal{T} then we may *flip* it to obtain a new triangulation \mathcal{T}' . This is done by replacing e with e' , the opposite diagonal of the square containing e , as shown in Figure 1.




Figure 1: Flipping an edge of a triangulation.

Proposition 2.5 ([9, Page 30]). *Let \mathcal{T} and \mathcal{T}' be triangulations of S and suppose that \mathcal{T}' is obtained from \mathcal{T} by flipping the edge e . Let a, b, c and d be the four edges of \mathcal{T} shown in Figure 1. Then for each multicurve $\gamma \in \mathcal{C}(S)$,*

$$\iota(\gamma, e') = \max(\iota(\gamma, a) + \iota(\gamma, c), \iota(\gamma, b) + \iota(\gamma, d)) - \iota(\gamma, e). \quad \square$$

We encapsulate these two moves in a simplicial *graph of triangulations* $G = G(S)$. The vertices of G correspond to triangulations of S . Two vertices in G are connected by an edge if and only if their triangulations differ by a reordering of their edges or by a single flip.

We assign each edge of this graph length one and denote the *length* of a path p in G with respect to the induced path metric by $\ell(p)$.

2.2 Encoding mapping classes

We note that G is connected [5, Page 190] and that the natural action of $\text{Mod}^+(S)$ on G is *geometric*, that is, a properly discontinuous, cocompact action by isometries. Therefore G and $\text{Mod}^+(S)$ are quasi-isometric [2, Proposition 8.19] and so we can use paths in G to combinatorially represent mapping classes. Specifically, we represent $h \in \text{Mod}^+(S)$ via a path p in G from \mathcal{T} to $h(\mathcal{T})$.

One method of obtaining such a path is to first fix paths representing each of the generators in X . Then we can construct a path p representing $h = h_1 \cdots h_k \in X^*$ by concatenating together translated copies of the paths representing h_1, \dots, h_k . Using this construction, $\ell(p) \in O(\ell(h))$

Now suppose that p is a path from \mathcal{T} to \mathcal{T}' . If we are given $v = \mathcal{T}(\gamma)$ then we may use Lemma 2.4 and Proposition 2.5 to compute $\mathcal{T}'(\gamma)$.

Lemma 2.6. *Suppose that p is a path from \mathcal{T} to \mathcal{T}' and that $\gamma \in \mathcal{C}(S)$ is a multicurve. If $v := \mathcal{T}(\gamma)$ is k -bounded then $v' := \mathcal{T}'(\gamma)$ is $(k + \ell(p))$ -bounded and can be computed in at most*

$$O(k\ell(p) + \ell(p)^2)$$

operations.

Proof. When p consists of a single relabelling or flip, Lemma 2.4 and Proposition 2.5 show that v' is $(k + 1)$ -bounded and can be computed in at most $O(k)$ operations. The result then holds by induction on $\ell(p)$. \square

In the case when p is a path representing $h \in \text{Mod}^+(S)$, as

$$\mathcal{T}(h(\gamma)) = h^{-1}(\mathcal{T})(\gamma),$$

this lemma allows us to compute $\mathcal{T}(h(\gamma))$ from $\mathcal{T}(\gamma)$ in $O(k\ell(p) + \ell(p)^2)$ operations.

Remark 2.7. By Alexander's trick, a mapping class fixes every edge vector if and only if it fixes every 1-bounded edge vector [3, Proposition 2.8]. However, if S is connected and not a once-marked torus or four times marked sphere then the only mapping class which fixes every edge vector is the identity map. Thus, by testing whether these vectors are fixed by h we can determine whether $h = \text{Id}$ in at most

$$O(\ell(h)^2)$$

operations. This gives the same asymptotic bound on the word problem for mapping class groups as Mosher obtained by showing that $\text{Mod}^+(S)$ is automatic [10, Section 3].

Now observe that, as the functions used in Lemma 2.4 and Proposition 2.5 are piecewise linear¹, for each pair of triangulations \mathcal{T} and \mathcal{T}' there is also a piecewise linear function $f: \mathbb{N}^\zeta \rightarrow \mathbb{N}^\zeta$ such that

$$\mathcal{T}'(\gamma) = f(\mathcal{T}(\gamma))$$

¹Here piecewise linear functions are required to have a finite number of cells, each defined by a system of linear inequalities.

for each multicurve $\gamma \in \mathcal{C}(S)$. We express this piecewise-linear function using two collections of matrices $\{A_i\}$ and $\{B_i\}$. Here the matrix A_i describes the linear transformation inside of the i^{th} cell of f while the matrix B_i describes the system of linear inequalities which define the i^{th} cell. Given a path p from \mathcal{T} to \mathcal{T}' we can construct these matrices as follows:

- If p consists of a single reordering then we define its matrices to be:

$$\begin{aligned} A_1 &:= \Sigma \\ B_1 &:= (0 \ \cdots \ 0) \end{aligned}$$

where Σ is the permutation matrix corresponding to the relabelling.

- If p consists of a single flip of an edge e of \mathcal{T} , as shown in Figure 1, then we define its matrices to be:

$$\begin{aligned} A_1 &:= \text{Id} + E_{ea} + E_{ec} - 2E_{ee} \\ A_2 &:= \text{Id} + E_{eb} + E_{ed} - 2E_{ee} \\ B_1 &:= E_a + E_c - E_b - E_d \\ B_2 &:= E_b + E_d - E_c - E_a \end{aligned}$$

where Id is the identity matrix, E_i is the $\zeta \times 1$ matrix with a 1 at position $(i, 1)$ and 0 everywhere else and E_{ij} is the $\zeta \times \zeta$ matrix with a 1 at position (i, j) and 0 everywhere else.

- Otherwise we decompose p as $p' \cdot p''$ and inductively define its matrices to be:

$$A_i := A_k'' \cdot A_j' \quad \text{and} \quad B_i := \left(\frac{B_j'}{B_k'' \cdot A_j'} \right)$$

where $\{A_j'\}$ and $\{B_j'\}$ are the matrices of p' and $\{A_k''\}$ and $\{B_k''\}$ are the matrices of p'' .

Lemma 2.8. *Suppose that p is a path from \mathcal{T} to \mathcal{T}' . Let $\{A_i\}$ and $\{B_i\}$ be the matrices defined above.*

1. Each A_i and B_i is $\ell(p)$ -bounded.
2. Each B_i has $O(\ell(p))$ rows.
3. For each multicurve $\gamma \in \mathcal{C}(S)$ we have that $B_i \cdot \mathcal{T}(\gamma) \geq 0$ for some i .
4. For each multicurve $\gamma \in \mathcal{C}(S)$, if $B_i \cdot \mathcal{T}(\gamma) \geq 0$ then $\mathcal{T}'(\gamma) = A_i \cdot \mathcal{T}(\gamma)$. \square

3 Determining and certifying reducibility

We will express the reducibility problem as an linear programming problem. Small invariant curves will then correspond to small solutions to this linear programming problem. This is closely related to the vertex enumeration problem for unbounded polytopes [1]. We start with a technical lemma for bounding determinants of matrices.

Lemma 3.1. *If M is a k -bounded, $n \times n$ matrix then $\det(M)$ is $(kn + n \log(n)/2)$ -bounded.*

Proof. This bound follows immediately from Hadamard's inequality [4, Theorem 14.1.1]. \square

Proposition 3.2. *Suppose that M is a k -bounded, $m \times n$ matrix. If the polytope*

$$P := \{v \in \mathbb{N}^n : M \cdot v \geq 0\}$$

is non-trivial then it contains a non-trivial $(nk + n \log(n)/2)$ -bounded integral vector.

Proof. Without loss of generality we may assume that the basis vectors E_i are rows of M . Let v_0 be an *extremal* vector of P , that is, $v_0 \in P - \{0\}$ and there are $n - 1$ linearly independent rows r_1, \dots, r_{n-1} of M such that $r_i \cdot v_0 = 0$. We claim that we can rescale v_0 to obtain $v_1 \in P - \{0\}$, a vector in which each entry is a $(nk + n \log(n)/2)$ -bounded integer.

To see this, define $r_0 := (1 \ \cdots \ 1)$ and let A be the matrix with rows r_0, r_1, \dots, r_{n-1} . Then v_0 is the unique solution to

$$A \cdot v = \|v_0\| \cdot \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}$$

By Cramer's rule, if A_i is the matrix obtained by replacing the i^{th} column of A by $(1 \ 0 \ \cdots \ 0)^T$ then the i^{th} entry of v_0 is given by

$$\|v_0\| \cdot \frac{\det(A_i)}{\det(A)}.$$

Hence, by rescaling v_0 by $|\det(A)|/|v_0|$ we obtain a vector $v_1 \in P - \{0\}$ whose i^{th} entry is $|\det(A_i)|$. However v_1 is $(nk + n \log(n)/2)$ -bounded as each $|\det(A_i)|$ is by Lemma 3.1. \square

Theorem 3.3. *Suppose that $h \in \text{Mod}^+(S)$ is a mapping class and that p is a path from \mathcal{T} to $h(\mathcal{T})$. If h is reducible then there is an h -invariant multicurve $\gamma \in \mathcal{C}(S)$ such that $\mathcal{T}(\gamma)$ is $O(\ell(p))$ -bounded.*

Proof. Let $\{A_i\}$ and $\{B_i\}$ be the matrices of Lemma 2.8. Additionally, let $\{F_j\}$ be the matrices of Lemma 2.3. Then for each i and j , let

$$M(i, j) := \begin{pmatrix} \frac{A_i - \text{Id}}{-(A_i - \text{Id})} \\ \frac{B_i}{F_j} \\ \text{Id} \end{pmatrix}$$

We begin by claiming that h is reducible if and only if there is a non-trivial solution to $M(i, j) \cdot v \geq 0$ for some i and j .

To prove this claim, firstly suppose that $h(\gamma) = \gamma$ and let $v := \mathcal{T}(\gamma) \neq 0$. Let i be such that $B_i \cdot v \geq 0$ and so $A_i \cdot v = \mathcal{T}(h(\gamma)) = v$. Hence

$$(A_i - \text{Id}) \cdot v \geq 0 \quad \text{and} \quad -(A_i - \text{Id}) \cdot v \geq 0.$$

As γ is a multicurve there is a j such that $F_j \cdot v \geq 0$. Thus v is a non-trivial vector and $M(i, j) \cdot v \geq 0$.

Conversely, suppose that v is a non-trivial solution to $M(i, j) \cdot v \geq 0$. Without loss of generality we may assume that the entries of v are non-negative and rational as $M(i, j)$ defines a rational polytope. Furthermore, by scaling v by a sufficiently large natural number we may assume that

$$F_j \cdot v \geq_2 0.$$

Hence, there is a multicurve $\gamma \in \mathcal{C}(S)$ such that $\mathcal{T}(\gamma) = v$. As $B_i \cdot v \geq 0$ and v lies in the kernel of $A_i - \text{Id}$, we have that $h(\gamma) = \gamma$. This proves the claim.

Now by Lemma 2.8 each A_i and B_i is $O(\ell(p))$ -bounded and so each $M(i, j)$ is too. Therefore by Proposition 3.2, there is a non-trivial $O(\ell(p))$ -bounded vector v_0 such that $M(i, j) \cdot v_0 \geq 0$. Then

$$F_j \cdot 2v_0 \geq_2 0.$$

Thus there is a multicurve $\gamma \in \mathcal{C}(S)$ such that $\mathcal{T}(\gamma) = 2v_0$. Hence it follows that γ is an h -invariant multicurve such that $\mathcal{T}(\gamma)$ is $O(\ell(p))$ -bounded. \square

As we may choose a path p from \mathcal{T} to $h(\mathcal{T})$ such that $\ell(p) \in O(\ell(h))$ we immediately obtain that:

Corollary 3.4. *Fix \mathcal{T} , a triangulation of S . If $h \in X^*$ is reducible then there is an h -invariant multicurve $\gamma \in \mathcal{C}(S)$ such that $\mathcal{T}(\gamma)$ is $O(\ell(h))$ -bounded.* \square

We may use such a multicurve as a certificate that $h \in X^*$ is reducible. Given its edge vector, using Lemma 2.3 we can first verify that it corresponds to a multicurve γ in $O(\ell(h))$ operations. Secondly, by using Lemma 2.6 we can compute $\mathcal{T}(h(\gamma))$ in $O(\ell(h)^2)$ time. Finally we can verify that $\mathcal{T}(h(\gamma)) = \mathcal{T}(\gamma)$, and so verify that h is reducible, in $O(\ell(h))$ time. This shows that:

Corollary 3.5. *The reducibility problem is in **NP**.* \square

As with all problems in **NP**, this also gives an exponential time algorithm to decide whether or not a mapping class is reducible. We iterate through the exponentially many $O(\ell(h))$ -bounded vectors; if there is a non-trivial one which corresponds to a multicurve and is fixed by h then h is reducible, if not then h is irreducible.

4 Subsurfaces

When $h \in \text{Mod}^+(S)$ is a reducible mapping class, as well as fixing a multicurve it also fixes a proper subsurface. In order to study the induced mapping class on such an invariant subsurface without talking about surfaces with boundary, we introduce the notion of *crushing* S along a multicurve $\gamma \in \mathcal{C}(S)$.

Definition 4.1. We *crush* S along γ to obtain the (again possibly disconnected) surface S_γ by:

1. removing an open regular neighbourhood of γ ,
2. collapsing the new boundary components to additional marked points, and then

3. removing any components that are twice marked spheres.

See Figure 2 for example.




Figure 2: Crushing along a multicurve.

Now if \mathcal{T} is a triangulation of S then we may track it as we crush S along a multicurve $\gamma \in \mathcal{C}(S)$. After possibly collapsing any bigons that are created, this results in a triangulation \mathcal{T}_γ of S_γ . There is a canonical bijection between the edges of \mathcal{T} and \mathcal{T}_γ and so $\zeta(S_\gamma) = \zeta$. To see this consider the following construction of \mathcal{T}_γ^* , the dual graph of \mathcal{T}_γ inside of S :

1. For each face $f \in F(\mathcal{T})$, place a vertex v in the *core* of f , that is, the component of $f - \gamma$ which meets all three sides of f .
2. Extend three half-edges from v to ∂f whilst avoiding γ .
3. Extend these half edges along the *corridors* created by parallel strands of γ until they connect with another half edge.

Proposition 4.2. *If p is a path from \mathcal{T} to \mathcal{T}' then crushing each triangulation of p along γ , and possibly discarding any repeated triangulations, gives a path p_γ from \mathcal{T}_γ to \mathcal{T}'_γ in $G(S_\gamma)$.*

Proof. The result clearly holds when p consists of a single reordering of the edges of \mathcal{T} . If p consists of a single flip then the combinatorics of \mathcal{T}_γ and \mathcal{T}'_γ agree away from the faces coming from the faces incident to the flipped edge. Thus \mathcal{T}_γ and \mathcal{T}'_γ share at least $\zeta - 1$ edges and so they are either equal or differ by a single flip. The result then follows for all paths by induction on $\ell(p)$. \square

In fact when \mathcal{T}' is obtained by flipping the edge e of \mathcal{T} , we have that \mathcal{T}_γ and \mathcal{T}'_γ are equal if and only if there is an arc of γ passing from one side of the square containing e to the opposite side. Following the notation of Figure 1, this occurs if and only if $\iota(\gamma, a) + \iota(\gamma, c) \neq \iota(\gamma, b) + \iota(\gamma, d)$.

Finally, we note that by construction $\ell(p_\gamma) \leq \ell(p)$.

4.1 Maximal curves

When $\gamma \in \mathcal{C}(S)$ is an h -invariant multicurve, we write $h_\gamma \in \text{Mod}^+(S_\gamma)$ for the mapping class induced on S_γ by h . Using this notation, if p is a path from \mathcal{T} to $h(\mathcal{T})$ then p_γ is a path from \mathcal{T}_γ to $h_\gamma(\mathcal{T}_\gamma)$.

Definition 4.3. A multicurve $\gamma \in \mathcal{C}(S)$ is *h -maximal* if it is h -invariant and h_γ is irreducible.

Now the bijection between edges of \mathcal{T} and the edges of \mathcal{T}_γ gives a map $\iota_\gamma: \mathcal{C}(S_\gamma) \rightarrow \mathcal{C}(S)$, lifting multicurves on S_γ back into S . Furthermore, if $\mathcal{T}(\gamma)$ is k -bounded then there is a k -bounded integer matrix M such that

$$\mathcal{T}(\iota_\gamma(\gamma')) = M \cdot \mathcal{T}_\gamma(\gamma').$$

However, it will be easier to work with the map:

$$\overline{\iota}_\gamma: \mathcal{C}(S_\gamma) \rightarrow \mathcal{C}(S) \quad \text{given by} \quad \overline{\iota}_\gamma(\gamma') := \iota_\gamma(\gamma') \cup \gamma.$$

Lemma 4.4. *Suppose that $\mathcal{T}(\gamma)$ is k -bounded. If $\gamma' \in \mathcal{C}(S_\gamma)$ is multicurve and $\mathcal{T}_\gamma(\gamma')$ is k' -bounded then $\mathcal{T}(\overline{\iota}_\gamma(\gamma'))$ is $(k + k' + \zeta)$ -bounded.* \square

We may repeat the construction of an invariant multicurve on S_γ and use this bound to control the complexity of the result we obtain back on S . To help us do this rigorously we introduce a second notion of complexity, closely related to the dimension of the curve complex of S_γ [8].

Definition 4.5. The *complexity* of a multicurve $\gamma \in \mathcal{C}(S)$ is

$$\xi(\gamma) := 3g(S_\gamma) + n(S_\gamma) - 3|S_\gamma|.$$

Now note that if $\gamma \in \mathcal{C}(S)$ and $\gamma' \in \mathcal{C}(S_\gamma)$ then

$$\xi(\overline{\iota}_\gamma(\gamma')) < \xi(\gamma).$$

Additionally, $\xi(\gamma) \leq \zeta$ and if $\xi(\gamma) = 0$ then $\mathcal{C}(S_\gamma) = \emptyset$ and so γ must be h -maximal.

Theorem 4.6. *Suppose that $h \in \text{Mod}^+(S)$ is a reducible mapping class and that p is a path from \mathcal{T} to $h(\mathcal{T})$. Then there is an h -maximal multicurve $\gamma \in \mathcal{C}(S)$ such that $\mathcal{T}(\gamma)$ is $O(\ell(p))$ -bounded.*

Proof. As h is reducible there is an h -invariant multicurve $\gamma \in \mathcal{C}(S)$ such that $\mathcal{T}(\gamma)$ is $O(\ell(p))$ -bounded by Theorem 3.3.

Now suppose that γ is not h -maximal. As h_γ is reducible, we can reapply Theorem 3.3 to the crushed path p_γ from \mathcal{T}_γ to $h_\gamma(\mathcal{T}_\gamma)$. As $\ell(p_\gamma) \leq \ell(p)$, we deduce that there is an h_γ -invariant multicurve $\gamma' \in \mathcal{C}(S_\gamma)$ such that $\mathcal{T}_\gamma(\gamma')$ is $O(\ell(p))$ -bounded.

Following this we redefine γ to be $\overline{\iota}_\gamma(\gamma')$. This is again an h -invariant multicurve and, by Lemma 4.4, is still $O(\ell(p))$ -bounded. However, doing this decreases $\xi(\gamma)$ and so after repeating this process at most ζ times γ must become h -maximal. \square

Again, as we may choose a path p from \mathcal{T} to $h(\mathcal{T})$ such that $\ell(p) \in O(\ell(h))$ we immediately obtain that:

Corollary 4.7. *Fix \mathcal{T} , triangulation of S . If $h \in X^*$ is reducible then there is an h -maximal multicurve $\gamma \in \mathcal{C}(S)$ such that $\mathcal{T}(\gamma)$ is $O(\ell(h))$ -bounded.* \square

4.2 The canonical curve system

The *canonical curve system* $\sigma(h) \in \mathcal{C}(S)$ of a mapping class $h \in \text{Mod}^+(S)$ is the intersection of all h -maximal multicurves [3, Page 373]. It is non-empty if and only if the mapping class is reducible and of infinite order [6, Theorem 4.44].

Koberda and Mangahas showed there is an exponential upper bound on the entries of $\mathcal{T}(\sigma(h))$ [7, Theorem 1]. Corollary 4.7 also provides an alternate proof of their theorem.

Proposition 4.8. *Suppose that $h \in \text{Mod}^+(S)$ is a mapping class and that p is a path from \mathcal{T} to $h(\mathcal{T})$. Then $\mathcal{T}(\sigma(h))$ is $O(\ell(p))$ -bounded.*

Proof. If $\sigma(h)$ is empty then the result holds trivially. Otherwise, h is reducible and so by Corollary 4.7 there is an h -maximal multicurve $\gamma \in \mathcal{C}(S)$ which is $O(\ell(h))$ -bounded. Therefore, as $\sigma(h) \subseteq \gamma$, we have that $\sigma(h)$ is $O(\ell(h))$ -bounded too. \square

Once more, as we may choose a path p from \mathcal{T} to $h(\mathcal{T})$ such that $\ell(p) \in O(\ell(h))$ we immediately obtain that:

Corollary 4.9 ([7, Theorem 1]). *Fix \mathcal{T} , a triangulation of S . For each word $h \in X^*$, the edge vector $\mathcal{T}(\sigma(h))$ of the canonical curve system of h is $O(\ell(h))$ -bounded.* \square

References

- [1] Endre Boros, Khaled Elbassioni, Vladimir Gurvich, and Kazuhisa Makino. Generating vertices of polyhedra and related problems of monotone generation. In *Polyhedral computation*, volume 48 of *CRM Proc. Lecture Notes*, pages 15–43. Amer. Math. Soc., Providence, RI, 2009.
- [2] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999.
- [3] Benson Farb and Dan Margalit. *A primer on mapping class groups*, volume 49 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 2012.
- [4] D. J. H. Garling. *Inequalities: a journey into linear analysis*. Cambridge University Press, Cambridge, 2007.
- [5] Allen Hatcher. On triangulations of surfaces. *Topology Appl.*, 40(2):189–194, 1991.
- [6] Yoshikata Kida. The mapping class group from the viewpoint of measure equivalence theory. *Mem. Amer. Math. Soc.*, 196(916):viii+190, 2008.
- [7] Thomas Koberda and Johanna Mangahas. An effective algebraic detection of the Nielsen-Thurston classification of mapping classes. *J. Topol. Anal.*, 7(1):1–21, 2015.

- [8] H. A. Masur and Y. N. Minsky. Geometry of the complex of curves. II. Hierarchical structure. *Geom. Funct. Anal.*, 10(4):902–974, 2000.
- [9] Lee Mosher. Tiling the projective foliation space of a punctured surface. *Trans. Amer. Math. Soc.*, 306(1):1–70, 1988.
- [10] Lee Mosher. Mapping class groups are automatic. *Ann. of Math. (2)*, 142(2):303–384, 1995.